

27

Bus subsystem **812** provides a mechanism for letting the various components and subsystems of computing device **810** communicate with each other as intended. Although bus subsystem **812** is shown schematically as a single bus, alternative implementations of the bus subsystem may use multiple busses.

Computing device **810** can be of varying types including a workstation, server, computing cluster, blade server, server farm, or any other data processing system or computing device. Due to the ever-changing nature of computers and networks, the description of computing device **810** depicted in FIG. **8** is intended only as a specific example for purposes of illustrating some implementations. Many other configurations of computing device **810** are possible having more or fewer components than the computing device depicted in FIG. **8**.

In situations in which certain implementations discussed herein may collect or use personal information about users (e.g., user data extracted from other electronic communications, information about a user's social network, a user's location, a user's time, a user's biometric information, and a user's activities and demographic information), users are provided with one or more opportunities to control whether information is collected, whether the personal information is stored, whether the personal information is used, and how the information is collected about the user, stored and used. That is, the systems and methods discussed herein collect, store and/or use user personal information only upon receiving explicit authorization from the relevant users to do so. For example, a user is provided with control over whether programs or features collect user information about that particular user or other users relevant to the program or feature. Each user for which personal information is to be collected is presented with one or more options to allow control over the information collection relevant to that user, to provide permission or authorization as to whether the information is collected and as to which portions of the information are to be collected. For example, users can be provided with one or more such control options over a communication network. In addition, certain data may be treated in one or more ways before it is stored or used so that personally identifiable information is removed. As one example, a user's identity may be treated so that no personally identifiable information can be determined. As another example, a user's geographic location may be generalized to a larger region so that the user's particular location cannot be determined.

While several implementations have been described and illustrated herein, a variety of other means and/or structures for performing the function and/or obtaining the results and/or one or more of the advantages described herein may be utilized, and each of such variations and/or modifications is deemed to be within the scope of the implementations described herein. More generally, all parameters, dimensions, materials, and configurations described herein are meant to be exemplary and that the actual parameters, dimensions, materials, and/or configurations will depend upon the specific application or applications for which the teachings is/are used. Those skilled in the art will recognize, or be able to ascertain using no more than routine experimentation, many equivalents to the specific implementations described herein. It is, therefore, to be understood that the foregoing implementations are presented by way of example only and that, within the scope of the appended claims and equivalents thereto, implementations may be practiced otherwise than as specifically described and claimed. Implementations of the present disclosure are directed to each

28

individual feature, system, article, material, kit, and/or method described herein. In addition, any combination of two or more such features, systems, articles, materials, kits, and/or methods, if such features, systems, articles, materials, kits, and/or methods are not mutually inconsistent, is included within the scope of the present disclosure.

The invention claimed is:

**1.** A method implemented by one or more processors, comprising:

receiving an initial voice input provided by a user via a client device;

performing a voice to text conversion to convert the initial voice input to initial text;

determining an intended action based on the initial text; identifying a mandatory parameter that is stored as mandatory for the intended action;

determining that the initial text lacks specification of any value for the mandatory parameter;

in response to determining that the initial text lacks specification of any value for the mandatory parameter: generating a natural language prompt based on the mandatory parameter, and

providing the natural language prompt as a reply to the initial voice input, the prompt being provided for presentation to the user via a user interface output device of the client device;

receiving additional voice input provided by the user in response to providing the natural language prompt;

determining a value for the mandatory parameter based on the additional voice input;

selecting a particular third-party agent from a group of third-party agents that can each perform the intended action;

transmitting a third-party invocation request that includes the value for the mandatory parameter, wherein the transmitting is to the particular third-party agent via one or more network interfaces;

receiving responsive content from the particular third-party agent in response to transmitting the intended action and the value, the receiving being via one or more of the network interfaces;

providing output that is based on the responsive content for presentation to the user;

receiving further voice input provided by the user in response to providing the output;

performing an additional voice to text conversion to convert the further voice input to further text;

transmitting the further text to the particular third-party agent;

in response to transmitting the further text to the particular third-party agent:

receiving further responsive content from the particular third-party agent, and

providing further output that is based on the further responsive content for presentation to the user;

receiving yet further voice input provided by the user in response to the further output;

determining that the yet further voice input indicates a desire to interact with another third-party agent; and

in response to determining that the yet further voice input indicates a desire to interact with another third-party agent:

transmitting, to an alternative third-party agent of the group of third party agents, an additional third-party invocation request that includes the value for the mandatory parameter.